International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Legal Insight: AI-Driven Bail Prediction System

Eva Shaji Zacharia[1], Warren Anthony Castelino[2], Aryan Puthran[3], Jayashri Mittal[4]

[1,2,3]Student, Department of Computer Engineering, St. Francis Institute of Technology
[4]Faculty, Department of Computer Engineering, St. Francis Institute of Technology

Abstract
The Legal Insight provides a digital platform to automate and accelerate the bail application process. By analysing key factors from case details and historical case data, the system helps undertrial prisoners and legal authorities determine bail eligibility quickly and fairly. This paper details our approach using IndicBERT, a language model pre-trained in Indian languages, which is fine-tuned on the HLDC dataset to handle bail prediction tasks. The project emphasises transparency, efficiency, and multi-language support in making fair predictions

Keywords: Bail Prediction, Explainable AI, IndicBERT, Indian Legal System, Legal Document Analysis, LegalTech

1. Introduction

The decision to grant bail is one of the most consequential in the criminal justice system. It affects not only the liberty of an individual awaiting trial but also the efficiency, credibility, and humaneness of the judicial system itself. In India, where over 75% of the prison population consists of undertrial prisoners, delays in bail decisions continue to contribute to overcrowding, prolonged detention, and unequal access to justice [1]. These delays are often a result of subjective human judgements, variations in legal interpretation, and systemic inefficiencies.

Amid these challenges, the emergence of Artificial Intelligence (AI) offers an opportunity to make legal decision-making more transparent, consistent, and accessible. AI tools are increasingly being explored for tasks ranging from case law analysis and legal summarisation to predictive modelling in sentencing and bail decisions. As Kumar highlights, machine learning algorithms such as XGBoost and SVM have shown promising results in forecasting judicial outcomes by analysing past rulings and legal patterns [2].

This trend reflects a larger shift toward data-driven justice, where predictive modelling is not intended to replace judges but to support them in making faster, fairer, and more consistent decisions. Generative AI is also beginning to influence how legal professionals learn and work. Prakash and Nair [3] argue that the integration of AI tools into legal education not only enhances research and analysis but also necessitates ethical training, especially as future lawyers will need to engage with AI systems critically and responsibly.

Moreover, NLP techniques have significantly advanced the way legal documents are handled. Sharma et al. [4] and Ghosh et al. [5] have demonstrated the importance of customised summarisation techniques to distil key information from lengthy legal texts, which is especially critical in Indian contexts where judgements are often verbose. Mahapatra et al. [6] further emphasise that legal access is not just a

technological issue but a linguistic one—highlighting the need for translation models that make legal text available in regional Indian languages.

Our research directly builds on this body of work. We propose an AI-based bail prediction system using IndicBERT, a multilingual transformer model fine-tuned on the Hindi Legal Documents Corpus (HLDC) introduced by Kapoor [7]. IndicBERT's architecture enables it to parse the linguistic and logical structures specific to Indian legal cases, including arguments, judicial summaries, and factual descriptions. This model is designed not merely to automate classification but to serve as an assistive system for judges, aligning with the "human-in-the-loop" philosophy advocated by Harasta et al. [8].

Further, we recognise the importance of legal precedent in prediction tasks. Wu et al. [9] have shown that large language models can be significantly enhanced when combined with domain-specific models that retrieve relevant precedents. These findings inform our post-processing stage, which integrates precedent alignment into bail outcome recommendations.

Finally, Siani [1] underscores a powerful motivation behind this work: the belief that "justice delayed is justice denied". In a system that continues to struggle with backlogs, delays, and inequalities, leveraging AI to expedite bail decisions—fairly and responsibly—is not just a technological innovation but a moral imperative.

2. Literature Review

A. Survey of existing systems:

The legal world is experiencing a silent yet profound transformation. Where once legal research was bound to books, clerks, and hours of manual reading, today, Large Language Models (LLMs) are beginning to reshape how lawyers, judges, and scholars interact with the law. These models—capable of reading, understanding, summarising, and even predicting legal text—mark a shift not just in technology, but in the way justice is administered.

As Shu et al. [10] demonstrate through their development of LawLLM, LLMs trained on domain-specific corpora can grasp legal nuance at a level previously unimaginable for machines. These models can cite precedent, analyse statutory language, and engage in rudimentary legal reasoning. Harasta et al. [8], however, caution that while LLM-generated documents are efficient, human judges and lawyers still prefer those they perceive to be authored by humans—raising ethical and professional questions about authorship and transparency.

Still, the value of LLMs is undeniable, especially in overburdened legal systems. In India, for instance, the staggering backlog of cases—over 4 crore pending at various levels—demands innovative solutions. This has led to growing interest in AI applications that go beyond information retrieval, stepping into areas like legal document summarisation, translation, and judgement prediction.

Initial work in this domain focused on summarisation, enabling professionals to grasp key ideas from complex and verbose judgements. Sharma et al. [4] provide a comprehensive analysis of summarisation techniques tailored for Indian legal documents, revealing the need for context-aware models that can capture legal arguments, precedents, and outcomes. Ghosh et al. [5] took a different route—emphasising text normalisation, particularly for Indian legal judgements that often contain archaic phrasing, code-mixed language, and unstructured formats.

But legal understanding is also deeply linguistic. Mahapatra et al. [6] argue that access to justice cannot be divorced from language accessibility. Their benchmark for translating legal texts into Indian languages underlines the challenges faced by non-English-speaking litigants, especially in regional and rural courts.

This work reinforces the need for multilingual models—such as IndicBERT—that understand Indian legalese beyond just English.

As these foundational tools matured, researchers turned toward prediction tasks—one of the most promising being Legal Judgement Prediction (LJP). Wu et al. [9] propose a hybrid approach where general-purpose LLMs collaborate with legal-domain models to make predictions based on not just facts, but legal precedent. Their Precedent-Enhanced LJP framework significantly improved the alignment of AI-generated outcomes with actual rulings, especially when legal context and history were integrated into model reasoning.

This progress, however, is still largely global in focus. There is a relative scarcity of India-specific predictive models, particularly in the bail domain, despite it being one of the most critical decision points in criminal trials.

India's criminal courts are overwhelmed. Over 75% of the prison population comprises undertrial prisoners—individuals who haven't been convicted but are awaiting bail decisions [1]. These decisions often hinge on subjective factors: the judge's perception of risk, the lawyer's skill, or even the timing of the hearing. In this context, AI could offer an objective, consistent, and explainable layer of decision support. Kapoor [7], in his thesis on Bail Prediction on Hindi Case Documents, takes an early step toward this goal. Using the Hindi Legal Documents Corpus (HLDC), he demonstrates that NLP models can indeed parse bail case files and produce prediction scores that correlate with judicial decisions. However, his model remains isolated—untested on newer multilingual architectures like IndicBERT and disconnected from larger LLM ecosystems.

Kumar [2] extends this idea to the broader realm of predictive modelling in legal decisions, showcasing how classifiers like SVM and Random Forest can be trained on structured case data. While promising, such models still require significant preprocessing and domain-specific feature engineering—a problem modern transformer-based LLMs are better equipped to handle things automatically.

Interestingly, the gap between legal education and AI deployment is also closing. Prakash and Nair [3] discuss how LLMs are now entering classrooms, teaching future lawyers not just case law, but how to think about law computationally. This reflects a cultural shift—one where the legal fraternity is beginning to see AI not as a threat but as a collaborator.

Yet, the practical deployment of such systems—especially for bail prediction in India—remains underdeveloped. Current models often lack transparency, ignore regional language data, or don't accommodate the nuanced reasoning found in Indian case law.

B. Limitations:

While recent advancements in legal AI have been promising, the literature reveals significant limitations that constrain the scalability and trustworthiness of current approaches. One recurring concern is the lack of explainability in judgement prediction systems. Models such as those discussed by Wu et al. [9] and Kapoor [7] primarily function as black boxes, offering little to no rationale behind their predictions, which is especially problematic in legal contexts where transparency is essential. Similarly, Harasta et al. [8] show that legal professionals remain sceptical of AI-generated documents, particularly when they cannot distinguish the reasoning behind outputs.

Furthermore, a strong bias toward Western legal systems is evident in several studies. For instance, LawLLM by Shu et al. [10] is trained specifically on U.S. legal data, while translation and summarisation models from Ghosh et al. [5] and Sharma et al. [4] either overlook Indian judicial nuances or restrict themselves to English-language texts. Although Mahapatra et al. [6] attempt to bridge this with Indian

language translation benchmarks, challenges in handling multilingual and unstructured legal corpora persist.

Data quality and structure also emerge as notable concerns. Studies like those by Kumar [2] and Siani [1] assume access to clean, labelled datasets and do not account for the heterogeneity and complexity of real court records. The lack of standard preprocessing strategies for Indian legal texts further hampers replicability and deployment. Additionally, Prakash and Nair [3] explore generative AI's role in education but do not extend their findings to high-stakes or legally consequential applications, leaving a gap between theoretical potential and practical integration.

Table 1. Summary of the Limitations

Citation	Limitations
[1] Siani (2024)	High-level exploration; lacks technical implementation and dataset detail
[2] Kumar (2024)	Assumes availability of clean, structured datasets; lacks data preparation strategies
[3] Prakash & Nair (2024)	Discusses GenAI in education only; no extension to real legal decision-making scenarios
[4] Sharma et al. (2023)	Summarization techniques not adapted for multilingual or domain-specific complexities in Indian law
[5] Ghosh et al. (2022)	Limited scope for summarisation; lacks integration with legal reasoning
[6] Mahapatra et al. (2023)	Focuses on translation; limited handling of unstructured or noisy Indian legal documents
[7] Kapoor (2022)	Bail prediction is opaque; no legal reasoning or explainability embedded
[8] Harasta et al. (2024)	Demonstrates user distrust in AI-authored legal content; lacks trust-building features
[9] Wu et al. (2023)	Judgment predictions lack explainability; model operates as a black box
[10] Shu et al. (2024)	Focused on U.S. legal system; lacks transferability to Indian legal contexts

C. Motivation:

In India, the process of determining whether an accused individual is eligible for bail is often delayed and inconsistent, primarily due to manual scrutiny of legal documents, limited access to legal expertise, and subjective judicial discretion. This not only slows down judicial proceedings but also disproportionately affects undertrial prisoners, especially those who lack the resources to secure timely legal assistance. Our project addresses this gap by building an AI-powered bail prediction system, aimed at bringing speed, consistency, and fairness to the bail decision-making process.

A major challenge in the legal system is the lack of early-stage clarity for clients and lawyers about the potential bailability of a case. Currently, this assessment relies heavily on manual document reading and interpretation, which is time-consuming and prone to error. By leveraging Large Language Models

(LLMs) to analyse FIRs and charge sheets, our system can predict the likelihood of bail, providing preliminary guidance to both legal professionals and defendants.

The legal process is also burdened by information asymmetry and unequal access to legal insight. Individuals in rural or underserved regions often have limited understanding of legal outcomes and procedures. Our system democratises this access by offering an automated, explainable prediction model, ensuring that users can receive timely, AI-assisted insights into their legal standing—regardless of their location or background.

Moreover, the current legal tech platforms in India are mostly directory-based or information-centric, lacking deep legal document analysis capabilities. Our project goes beyond surface-level assistance by incorporating IPC tagging and AI-based document reasoning, making the prediction process not only data-driven but also grounded in the statutory framework.

Finally, issues of judicial backlog and procedural delays remain persistent in the Indian judicial system. Automating bail predictions can serve as an auxiliary tool for legal practitioners, helping them prepare stronger arguments and manage cases more effectively, while potentially reducing unnecessary detentions and improving the overall efficiency of the legal system.

D. Figures

Figure 1: Architecture

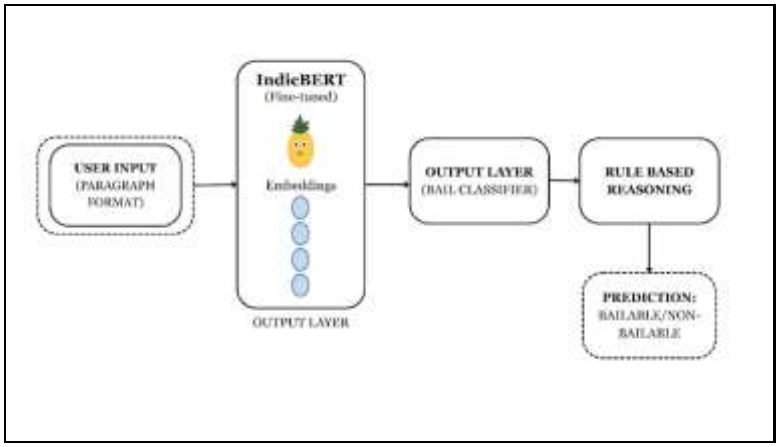

Figure 2: Use Case of Proposed System

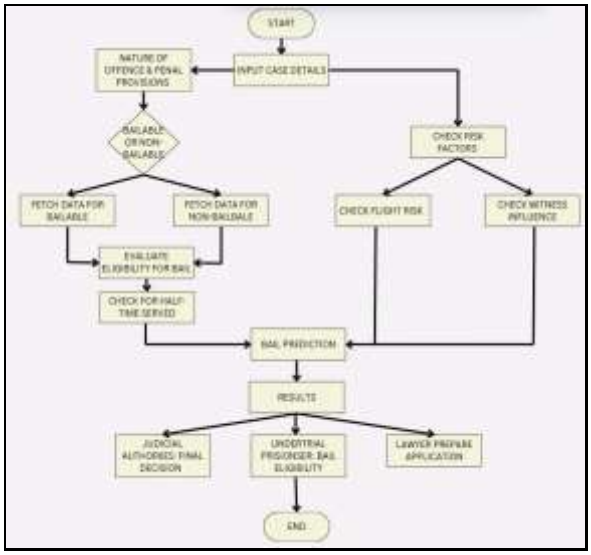

3. Proposed Methodology

The proposed system is an AI-based bail prediction model designed to determine whether a case is bailable or non-bailable based solely on legal document analysis. At its core, the system utilises IndicBERT, a transformer-based language model pre-trained on Indian languages, which enables it to process and understand legal documents such as FIRs and charge sheets written in regional languages.

The model is fine-tuned on the HLDC (Hindi Legal Documents Corpus), which contains annotated legal case data. This ensures the predictions are grounded in real-world legal precedents and language patterns used in actual bail hearings. The prediction task is treated as a binary classification problem, where the output is either "Bailable" or "Non-Bailable".

Key features of the methodology include:

1. IPC-based Legal Reasoning Layer: After the initial prediction, a post-processing logic layer maps case content to relevant Indian Penal Code (IPC) sections and verifies whether the predicted classification aligns with statutory definitions of bailable and non-bailable offences. This enhances reliability and legal alignment without requiring explainability from the model itself.
2. Multilingual Text Processing: IndicBERT's language capabilities allow the model to understand documents in Hindi and other Indian languages, which is critical for accessibility and adoption across diverse regions.
3. Evaluation Metrics: The model is evaluated using standard classification metrics such as accuracy, precision, recall, and F1 score (including macro and weighted averages). These metrics help assess the robustness and fairness of the predictions, especially for imbalanced datasets.
4. Survey with Legal Experts: To validate the real-world utility of the system, feedback is collected from lawyers and legal professionals through structured interviews or surveys. This helps in refining the model based on their input on practical relevance, perceived accuracy, and concerns such as over-reliance on AI for sensitive legal decisions.

4. Result And Evaluation

A. Results of the system:

To evaluate the performance of the proposed bail prediction system, standard classification metrics were employed—namely, precision, recall, F1-score, accuracy, and AUC (Area Under the ROC Curve). These metrics offer a comprehensive assessment of the model's ability to differentiate between bailable and non-bailable cases.

- Precision reflects the proportion of correctly predicted bailable (positive) cases among all predicted bailable cases.
- Recall indicates the proportion of actual bailable cases correctly identified.
- The F1-score, the harmonic mean of precision and recall, balances both metrics and is suitable for imbalanced legal datasets.
- Accuracy measures the overall percentage of correct predictions.
- AUC, derived from the Receiver Operating Characteristic (ROC) curve, quantifies the model's capability to distinguish between classes, with higher values indicating better discrimination.
- Evaluation was conducted on the HLDC test set comprising 35,400 labelled legal case records.

Table 2 presents the detailed classification report. The model achieved an overall accuracy of 82%, with particularly strong performance in predicting bailable cases (F1-score: 0.85) as compared to non-bailable cases (F1-score: 0.76). The macro average F1-score of 0.80 highlights balanced predictive performance,

while the weighted average F1-score of 0.82 suggests robustness to class distribution. Additionally, the AUC score of 0.88 confirms high discriminative ability, reinforcing the model's suitability for legal decision-support applications.

Table 2. Classification Report

Classification Report: Accuracy: 82%				
	precision	recall	F1 score	support
Non-bailable	0.76	0.75	0.76	13423
Bailable	0.85	0.85	0.85	21977
accuracy			0.82	35400
macro avg	0.80	0.80	0.80	35400
weighted avg	0.81	0.82	0.82	35400

To further assess the model's discriminative capability, the Receiver Operating Characteristic (ROC) curve was plotted. This curve illustrates the trade-off between the true positive rate and false positive rate across various threshold settings. The Area Under the Curve (AUC) score of 0.88 confirms the model's high effectiveness in distinguishing between bailable and non-bailable cases.

Figure 3: ROC Curve for Bail Prediction

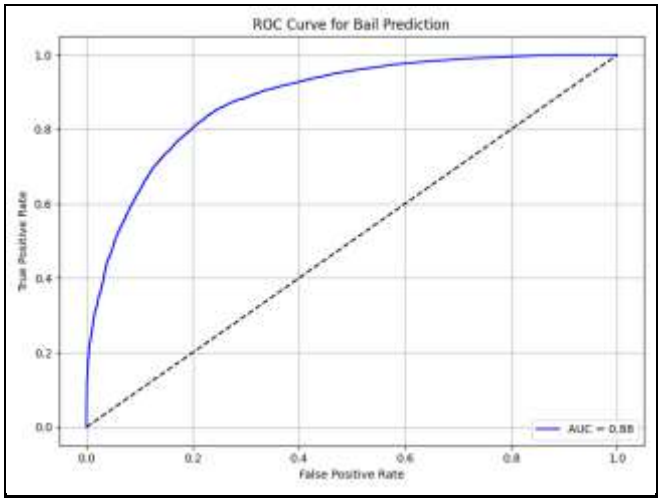

B. Survey from Lawyers and Law Firms:

To evaluate the practicality and effectiveness of the bail prediction system, a survey was conducted using Google Forms, targeting lawyers and legal professionals. The survey gathered responses on the system's accuracy, usability, efficiency, and alignment with legal reasoning, along with open-ended feedback on limitations and improvements. The responses were analysed to assess the system's relevance in real-world legal practice.

Table 3. Key Findings from survey

SR. NO.	Metric	Key Findings
1	Accuracy of Bail Predictions	Participants were asked how accurately the system predicts bail outcomes. The majority of respondents (75%) rated the system as very accurate.
2	Usability of the System	The system's ease of use was assessed, with 100% of respondents finding it very easy to use
3	Time Efficiency Compared to Manual Analysis	When asked about time savings, 87.5% of respondents stated that the system saves a lot of time, while 12.5% felt it saves some time
4	System Limitations	While some respondents reported no major issues, key concerns included the system's reliance on historical data, which may lead to errors if legal provisions change, and its potential struggle with unseen cases. Ethical concerns were also raised regarding AI-driven legal decisions. Additionally, a suggestion was made to improve accuracy by focusing on key languages like English, Hindi, and Marathi rather than relying solely on multilingual training. These insights highlight the need for better adaptability to legal updates and enhanced language-specific performance.
5	Ways to Improve the System	Respondents suggested expanding the system beyond criminal law to include other legal domains, ensuring regular updates with new case laws and acts. Maintaining high-quality data, especially when incorporating cases from different states, was emphasised. Security enhancements were also recommended for future development. Additionally, it was suggested that once the system is well-established, more features could be added, and a patent could be pursued.
6	Recommend this system for legal practice.	100% of the participants would recommend this.

Figure 4: Accuracy Of Bail Prediction

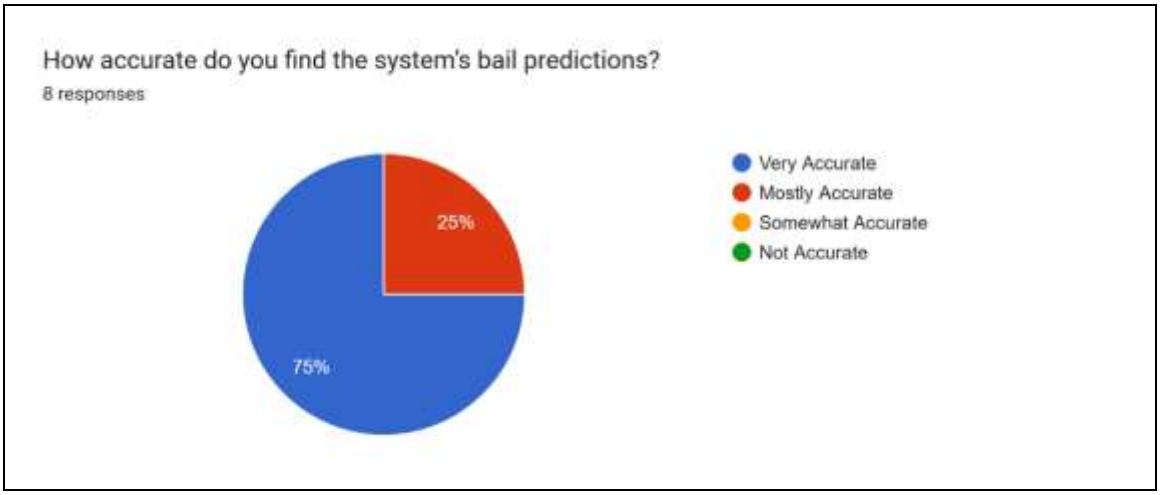

Figure 5: Ease of using The System

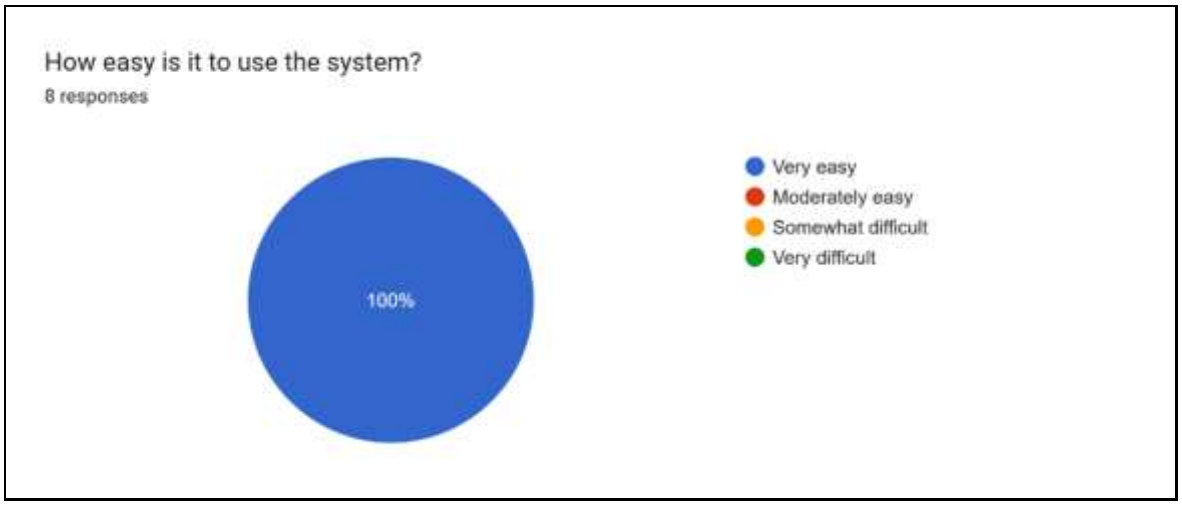

Figure 6: Time Efficiency Compared to Manual Analysis

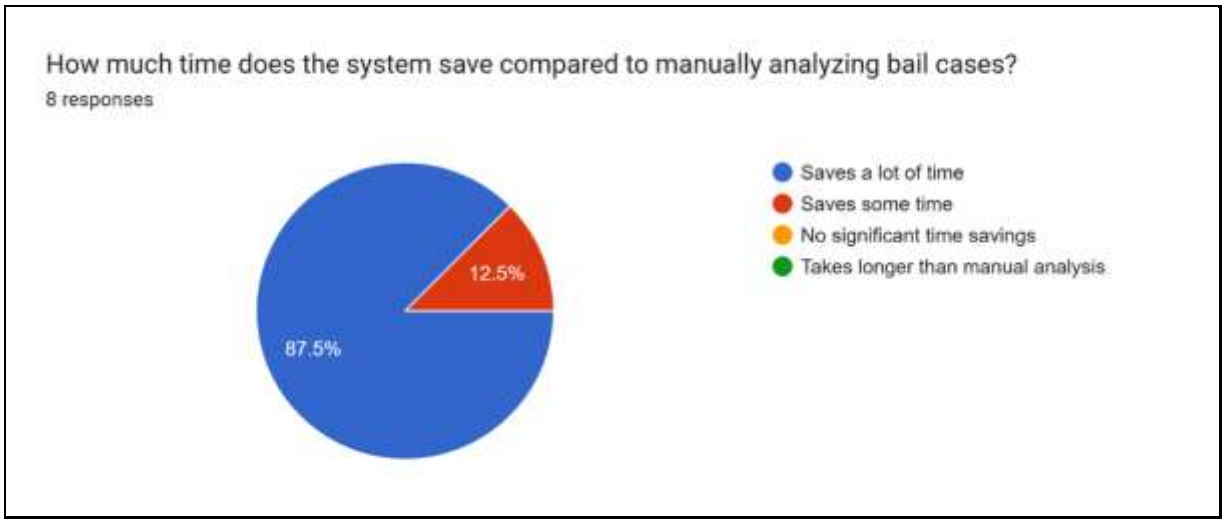

Figure 7: Recommend This System for Legal Practice

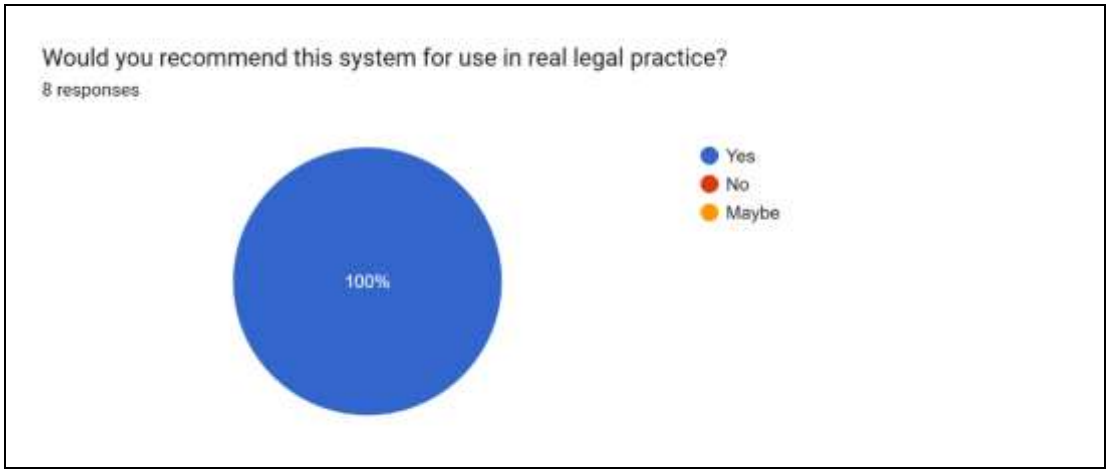

5. Conclusion

The proposed Bail Prediction System, built using IndicBERT and evaluated on the HLDC dataset, demonstrates strong potential in automating the assessment of bail eligibility for undertrial prisoners. With an AUC of 0.88 and an overall accuracy of 82%, the model offers a promising foundation for legal AI applications in India. However, to fully harness its capabilities, further enhancements are needed. Future work should focus on improving the model's ability to interpret the nuanced structure and terminology of legal texts, particularly those involving complex case narratives. Additionally, expanding linguistic support beyond Hindi is crucial for deployment across India's diverse legal landscape. Addressing challenges related to unstructured legal data and integrating more robust precedent-based reasoning will further increase the system's reliability and practical value in real-world judicial scenarios.

6. Future Work

While the current multilingual Bail Prediction System has shown promising results, several areas offer opportunities for further improvement and expansion. One significant area is the incorporation of larger and more diverse datasets. By integrating case data from various jurisdictions and legal contexts, the model could be better generalised, improving its applicability across different legal systems. Additionally, expanding the system's multilingual support to include more regional languages would enhance accessibility, ensuring that users from diverse linguistic backgrounds can benefit from its features. Real-time case data integration also stands as a crucial future development. Incorporating dynamic updates from courts and legal documents would enhance the system's accuracy and timeliness in predicting bail eligibility, offering more relevant case insights. Lastly, addressing model explainability and transparency is vital for improving user trust. By providing clear visualisations or breakdowns of how predictions are made, the system can ensure that legal professionals can easily interpret its results, ultimately aiding in better-informed decision-making.

7. References

1. Siani, J. A. "Empowering justice: Exploring the applicability of AI in the judicial system." Journal of Law and Legal Research Development (2024): 24-28.
2. Kumar, Mukesh. "Predictive Modelling in Legal Decision-Making: Leveraging Machine Learning for Forecasting Legal Outcomes." J. Electrical Systems 20, no. 3 (2024): 2060-2071.

3. Aswathy Prakash, G., and Vishnu Nair. "Integrating Generative AI into Legal Education: From Casebooks to Code, Opportunities and Challenges." Law, Technology and Humans 6, no. 3 (2024).
4. Sharma, S., Srivastava, S., Verma, P., Verma, A., & Chaurasia, S. N. (2023). A Comprehensive Analysis of Indian Legal Documents Summarization Techniques. SN Computer Science, 4(614).
5. Ghosh, S., Dutta, M., & Das, T. (2022). Indian Legal Text Summarization: A Text Normalization-based Approach. arXiv preprint arXiv:2206.06238v2
6. Mahapatra, S., Datta, D., Soni, S., Goswami, A., & Ghosh, S. (2023). Improving Access to Justice for the Indian Population: A Benchmark for Evaluating Translation of Legal Text to Indian Languages. arXiv preprint arXiv:2310.09765v1.
7. Kapoor, Arnav. "Justice Delayed is Justice Denied: Enabling Legal Artificial Intelligence via Bail Prediction on Hindi Case Documents." PhD diss., International Institute of Information Technology Hyderabad, 2022.
8. Harasta, J., Novotná, T., & Savelka, J. (2024). It Cannot Be Right If It Was Written by AI: On Lawyers' Preferences of Documents Perceived as Authored by an LLM vs a Human. arXiv preprint arXiv:2407.06798v1.
9. Wu, Y., Zhou, S., Liu, Y., Lu, W., Liu, X., Zhang, Y., Sun, C., Wu, F., & Kuang, K. (2023). Precedent-Enhanced Legal Judgment Prediction with LLM and Domain-Model Collaboration. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, pages 12060–12075. Association for Computational Linguistics
10. Shu, D., Zhao, H., Liu, X., Demeter, D., Du, M., & Zhang, Y. (2024). LawLLM: Law Large Language Model for the US Legal System. Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24), October 21-25, 2024, Boise, ID, USA. ACM.
11. Su, Weihang, et al. "JuDGE: Benchmarking Judgment Document Generation for Chinese Legal System." arXiv preprint arXiv:2503.14258 (2025).
12. Koenecke, Allison, et al. "Tasks and Roles in Legal AI: Data Curation, Annotation, and Verification." arXiv preprint arXiv:2504.01349 (2025).
13. Liang, Yueqing, et al. "Benchmarking LLMs for Political Science: A United Nations Perspective." arXiv preprint arXiv:2502.14122 (2025).
14. Rath, Alisha. AI Powered Legal Decision Support System. Diss. San José State University, 2025.
15. Branting, L. Karl, et al. "Scalable and explainable legal prediction." Artificial Intelligence and Law 29 (2021): 213-238.
16. Ejjami, Rachid. "AI-driven justice: Evaluating the impact of artificial intelligence on legal systems." Int. J. Multidiscip. Res 6.3 (2024): 1-29.
17. Tsihrintzis, George A., et al. Advances in Artificial Intelligence-Empowered Decision Support Systems. 2024.
18. Chalkidis, Ilias, Ion Androutsopoulos, and Nikolaos Aletras. "Neural legal judgment prediction in English." arXiv preprint arXiv:1906.02059 (2019).
19. Cao, Lang, et al. "PILOT: legal case outcome prediction with case law." arXiv preprint arXiv:2401.15770 (2024).
20. Araújo, Gabriele S., et al. "The Artificial Intelligence Integration in the Brazilian Legal Sector: A Systematic Review." Simpósio Brasileiro de Sistemas de Informação (SBSI) (2025): 575-584